arXiv:1605.05671v1 [math.ST] 18 May 2016

Sub-optimality of some continuous shrinkage priors

Anirban Bhattacharya, David B. Dunson, Debdeep Pati, Natesh S. Pillai

Texas A& M University, Duke University, Florida State University, Harvard University

Abstract

Two-component mixture priors provide a traditional way to induce sparsity in high-dimensional Bayes models. However, several aspects of such a prior, including computational complexities in high-dimensions, interpretation of exact zeros and non-sparse posterior summaries under standard loss functions, has motivated an amazing variety of continuous shrinkage priors, which can be expressed as global-local scale mixtures of Gaussians. Interestingly, we demonstrate that many commonly used shrinkage priors, including the Bayesian Lasso, do not have adequate posterior concentration in high-dimensional settings.

Keywords: Bayesian, Convergence rate, High dimensional, Lasso, ℓ_1, Lower bound, Penalized regression, Regularization, Shrinkage prior, sub-optimal

1. Introduction

With the recent flurry of activities in high-throughput data, taking advantage of sparsity to perform statistical inference is a common theme in situations where the number of model parameters (p) increases with the sample-size (n). In such scenarios, penalization methods [1] can yield a point estimate very quickly. There is a rich theoretical literature justifying the optimality properties of such penalization approaches [2, 3, 4, 5, 6, 7], with fast algorithms [8] and compelling applied results leading to their routine use.

On the other hand, statistical theory for characterizing the uncertainty of model parameters using penalization methods in high dimensions has received comparatively less attention [9]. Bayesian approaches provide a natural measure of uncertainty through the induced posterior distribution. Most penalization methods have a Bayesian counterpart. For instance, ℓ_1 and ℓ_2 regularization methods are equivalent to placing zero-mean double-exponential and Gaussian priors respectively on the parameter vector and

the solutions of the corresponding optimization problems are precisely the mode of the Bayesian posterior distribution. Moreover, a Bayesian approach has distinct advantages in terms of choice of tuning parameters, allowing key penalty parameters to be marginalized over the posterior distribution instead of relying on cross-validation. Thus a fruitful line of research is to investigate the behavior of the entire posterior distribution of the Bayesian models corresponding to penalization methods.

The process of eliciting prior distributions can be very tricky in high-dimensions. Two-component mixture priors with a point mass at zero are traditionally used in high-dimensional settings because of their ability to produce exact zeros and ease of eliciting hyperparameters based on the prior knowledge about the level of sparsity and the size of the signal coefficients. In [10, 11], the authors showed optimality properties for carefully chosen point mass mixture priors in high-dimensional settings. Recently, in an insightful article [12], several arguments were raised against the point mass priors concerning interpretation of exact zeros and computational issues arising from exploring a very high-dimensional model space. This prompted the authors of [12] to seek for continuous analogues of point mass priors based on Gaussian scale mixtures which obviates the need to search over the huge model space. These scale mixtures of Gaussian priors are designed to have a sharp peak near zero with heavy tails so as to emulate the point mass mixture priors. In the last few years, a huge variety of shrinkage priors have been proposed in the Bayesian literature [13, 14, 15, 16, 17]. In [17] the authors studied shrinkage priors and provided simple sufficient conditions for posterior consistency in $p \leq n$ settings. However, results on quantifying posterior concentration using continuous shrinkage priors are scarce.

Even from a purely practical point of view, considerable difficulties have arisen when attempts have been made to reflect prior beliefs on sparsity through the associated hyperparameters of these distributions. For example, suppose we wish to estimate $\theta_0 \in \mathbb{R}^n$ from $y \sim \mathrm{N}_n(\theta_0, I_n)$ under the prior knowledge that only a fraction of the coordinates of θ_0 are non-zero. What are the appropriate parameters one should choose in the Bayes Lasso formulation [13] to ensure efficient estimation of θ_0? A first step towards answering such questions is to understand the concentration of shrinkage priors around sparse vectors. This is critically important in two aspects. First, optimal prior concentration is almost necessary for optimal posterior contraction rates under a variety of loss functions. Second, studying the concentration of shrinkage priors around sparse vectors will yield insights into the geometry

of shrinkage priors which can then be harnessed for prior elicitation for a broad class of models.

Our contribution in this paper is two fold. First, we obtain sharp bounds for the concentration of continuous shrinkage priors around sparse high dimensional vectors. This is quite challenging because the joint distributions of such priors obtained through integrating several latent hyperparameters are often unwieldy to work with. One of the reasons why the point mass priors enjoy theoretical optimality properties is because they have optimal concentration around sparse vectors [10, 11]. We show that the concentration of some of the commonly used continuous shrinkage priors can sometimes be smaller than that of the point mass priors by several orders of magnitude. Second, using these results, we show that for the normal means problem, the Bayesian Lasso [13] has sub-optimal posterior concentration around sparse vectors under the ℓ_2 loss. Although not rigorously proven, we suspect the sub-optimality to be actually by a genuine power of the sample size, making this a serious concern against routine use of such priors. The negative results about shrinkage priors obtained in this paper contribute further to our understanding of the 'statistical efficiency *vs.* computational efficiency' trade off in high dimensional estimation. Technically, we rely on exploiting the conditional Gaussian formulation of continuous shrinkage priors and use results for small-ball probability of Gaussian distributions. Although we focus exclusively on the normal means problem, the ideas in this paper are applicable to other models including linear regression and to estimating high dimensional covariance matrices using Gaussian latent factor models [18].

One of the key insights emerging from this paper is that to induce an appropriate dependence structure in the scale parameters to ensure that the concentration is sufficiently high. In a recent article [19], we use this insight to construct a prior distribution on high dimensional vectors called the 'Dirichlet-Laplace' prior which enjoys optimal posterior concentration rates. Also, while completing this paper, we became aware of a recent article [20] in which the authors establish optimal posterior concentration for horseshoe priors in the normal means problem. Their techniques rely on explicit expressions for the posterior means and variances in the normal means model. However, the hyperparameters of the horseshoe priors are chosen using an empirical Bayes approach. The optimality of the fully Bayesian version of the horseshoe prior as in [16] remains an interesting open problem.

Proofs of technical results are provided in Section 6. Proofs of auxiliary lemmata stated in Section 5 are in a supporting document.

2. Preliminaries

Given sequences a_n, b_n, we denote $a_n = O(b_n)$ if there exists a global constant C such that $a_n \leq Cb_n$ and $a_n = o(b_n)$ if $a_n/b_n \to 0$ as $n \to \infty$. For a vector $x \in \mathbb{R}^r$, $\|x\|_2$ denotes its Euclidean norm. We will use Δ^{r-1} to denote the $(r-1)$-dimensional simplex $\{x = (x_1, \ldots, x_r)^{\mathrm{T}} : x_j \geq 0, \sum_{j=1}^r x_j = 1\}$. Further, let Δ_0^{r-1} denote $\{x = (x_1, \ldots, x_{r-1})^{\mathrm{T}} : x_j \geq 0, \sum_{j=1}^{r-1} x_j \leq 1\}$.

For a subset $S \subset \{1, \ldots, n\}$, let $|S|$ denote the cardinality of S and define $\theta_S = (\theta_j : j \in S)$ for a vector $\theta \in \mathbb{R}^n$. Denote $\mathrm{supp}(\theta)$ to be the *support* of θ, the subset of $\{1, \ldots, n\}$ corresponding to the non-zero entries of θ. Let $\ell_0[q; n]$ denote the subset of $\mathbb{R}^n$ consisting of q-sparse vectors θ with $|\mathrm{supp}(\theta)| \leq q$:

$$\ell_0[q; n] = \{\theta \in \mathbb{R}^n \ : \ \#(1 \leq j \leq n : \theta_j \neq 0) \leq q\}.$$

Let $\mathrm{DE}(\tau)$ denote a zero mean double-exponential or Laplace distribution with density $f(y) = (2\tau)^{-1} e^{-|y|/\tau}$ for $y \in \mathbb{R}$.

3. Concentration properties of global-local priors

3.1. Motivation

For a high-dimensional vector $\theta \in \mathbb{R}^n$, a natural way to incorporate sparsity in a Bayesian framework is to use point mass mixture priors

$$\theta_j \sim (1 - \pi)\delta_0 + \pi g_\theta, \quad j = 1, \ldots, n, \tag{1}$$

where $\pi = \Pr(\theta_j \neq 0)$, $\mathbb{E}\{|\mathrm{supp}(\theta)| \mid \pi\} = n\pi$ is the prior guess on model size (sparsity level), and g_θ is an absolutely continuous density on $\mathbb{R}$. A beta prior on π leads to an automatic multiplicity adjustment [21]. In [10], the authors established that prior (1) with an appropriate beta prior on π and suitable tail conditions on g_θ leads to a frequentist minimax optimal rate of posterior contraction in the normal means setting. We shall revisit the normal means problem in Section 4.

As mentioned before, the authors in [12] noted certain unappealing aspects of (1) including computational complexities in high-dimensions, interpretation of exact zeros and non-sparse posterior summaries under common loss functions. This has motivated a variety of continuous shrinkage priors [13, 16, 15, 22, 17], resembling the two-component priors facilitating

computation and interpretability. Almost all such shrinkage priors can be represented as global-local (GL) mixtures of Gaussians [12],

$$\theta_j \sim \mathrm{N}(0, \psi_j \tau), \quad \psi_j \sim f, \quad \tau \sim g, \tag{2}$$

where τ controls global shrinkage towards the origin while the local scales $\{\psi_j\}$ allow deviations in the degree of shrinkage. If g puts sufficient mass near zero and f is appropriately chosen, the hope is that GL priors in (2) can serve as a good enough proxy for (1) through a continuous density concentrated near zero with heavy tails.

The normal scale mixture representation in (2) allows for conjugate updating of θ and ψ in a block facilitating computation in high dimensions. Moreover, a number of frequentist regularization procedures such as ridge regression, lasso, bridge and elastic net correspond to posterior modes under GL priors with appropriate choices of f and g. For example, one obtains a double-exponential prior corresponding to the popular ℓ_1 or lasso penalty if f has an exponential distribution. However, unlike variable selection priors (1), many aspects of shrinkage priors are poorly understood. For example, even basic properties, such as the appropriate choices of f and g to have adequate prior concentration around a sparse vector are unknown till date.

There has been a recent awareness of these issues, motivating a basic assessment of the marginal properties of shrinkage priors for a single θ_j. Recent priors such as the horseshoe [16] and generalized double Pareto [17] are carefully formulated to obtain marginals having a high concentration around zero with heavy tails. This is well justified, but as we will see below, such marginal behavior alone is not sufficient; it is necessary to study the joint distribution of θ on $\mathbb{R}^n$. Specifically, we recommend studying the prior concentration $\mathbb{P}(\|\theta - \theta_0\| < t_n)$ where the true parameter θ_0 is assumed to be sparse: $\theta_0 \in \ell_0[q_n; n]$ with the number of non-zero components $q_n \ll n$ and

$$t_n = n^{\delta/2} \quad \text{with} \quad \delta \in (0, 1). \tag{3}$$

In models where $q_n \ll n$, the prior must place sufficient mass around sparse vectors to allow for good posterior contraction; see Section 4 for further details. Now, as a first illustration, consider the following two extreme scenarios: i.i.d. standard normal priors for the individual components θ_j *vs.* point mass mixture priors given by (1).

Theorem 3.1. *Assume that $\theta_0 \in \ell_0[q_n; n]$ with $q_n = o(n)$. Then, for i.i.d standard normal priors on θ_j,*

$$\mathbb{P}(\|\theta - \theta_0\|_2 < t_n) \leq e^{-cn}. \tag{4}$$

For point mass mixture priors (1) with $\pi \sim Beta(1, n+1)$ and g_θ being a standard Laplace distribution $g_\theta \equiv \mathrm{DE}(1)$,

$$\mathbb{P}(\|\theta - \theta_0\|_2 < t_n) \geq e^{-c \max\{q_n, \|\theta_0\|_1\}}. \tag{5}$$

Proof. Using $\|\theta\|_2^2 \sim \chi_n^2$, the claim made in (4) follows from an application of a Gaussian small probability in Lemma 5.1 and standard chi-square deviation inequalities. In particular, the exponentially small concentration also holds for $\mathbb{P}(\|\theta_0\|_2 < t_n)$. The second claim (5) follows from results in [10]. □

As seen from Theorem 3.1, the point mass mixture priors have much improved concentration around sparse vectors, as compared to the i.i.d. normal prior distributions. The theoretical properties enjoyed by the point mass mixture priors can partly be attributed to this improved concentration. The above comparison suggests that it is of merit to evaluate a shrinkage prior in high dimensional models under sparsity assumption by obtaining its concentration rates around sparse vectors. The main results are given in Section 6. Recall the GL priors presented in (2) and the sequence t_n in (3).

3.2. Prior concentration for global priors

This simplified setting involves only a global parameter, *i.e.*, $\psi_j = 1$ for all j. This subclass includes the important example of ridge regression, with τ routinely assigned an inverse-gamma prior, $\tau \sim \mathrm{IG}(\alpha, \beta)$.

Theorem 3.2. *Assume $\theta \sim$ GL with $\psi_j = 1$ for all j. If the prior g on the global parameter τ has an $\mathrm{IG}(\alpha, \beta)$ distribution, then*

$$\mathbb{P}(\|\theta\|_2 < t_n) \leq e^{-Cn^{1-\delta}}, \tag{6}$$

where $C > 0$ is a constant depending only on α and β and δ is from (3).

The above theorem shows that compared to i.i.d. normal priors (4), the prior concentration does not improve much under an inverse-gamma prior on the global variance regardless of the hyperparameters (provided they don't scale with n) even when $\theta_0 = 0$. Concentration around θ_0 away from zero

will clearly be even worse. Hence, such a prior is not well-suited in high-dimensional settings, confirming empirical observations documented in [23, 24]. It is also immediate that the same concentration bound in (6) would be obtained for the giG family of priors on τ.

In [24], the authors instead recommended a half-Cauchy prior as a default choice for the global variance (also see [23]). We consider the following general class of densities on $(0, \infty)$ for τ, to be denoted $\mathcal{G}$ henceforth, that satisfy: (i) $g(\tau) \leq M$ for all $\tau \in (0, \infty)$ (ii) $g(\tau) > 1/M$ for all $\tau \in (0, 1)$, for some constant $M > 0$. Clearly, $\mathcal{G}$ contains the half-Cauchy and exponential families. The following result provides concentration bounds for these priors.

Theorem 3.3. *Let $\|\theta_0\|_2 = o(\sqrt{n})$. Recall t_n and δ from (3). If the prior g on the global parameter τ belongs to the class $\mathcal{G}$ above then,*

$$C_1 e^{-(1-\delta)\log n} \leq \mathbb{P}(\|\theta\|_2 < t_n) \leq C_2 e^{-(1-\delta)\log n}. \tag{7}$$

Furthermore, if $\frac{1}{4}\|\theta_0\|_2 > t_n$, then

$$e^{-c_1 n \log a_n} \leq \mathbb{P}(\|\theta - \theta_0\|_2 < t_n) \leq e^{-c_2 n \log a_n}, \tag{8}$$

where $a_n = \|\theta_0\|_2 / \sqrt{2} t_n > 1$ and $c_i, C_i > 0$ are constants with C_1, C_2 depending only on M in the definition of $\mathcal{G}$ and c_1 depending on M and δ.

Equation (7) in Theorem 3.3 shows that the prior concentration around zero can be dramatically improved from exponential to polynomial with a careful prior on τ that can assign sufficient mass near zero, such as the half-Cauchy prior [23, 24]. Unfortunately, as (8) shows, for signals of large magnitude one again obtains an exponentially decaying probability. Hence, Theorem 3.3 conclusively shows that global shrinkage priors are simply not flexible enough for high-dimensional problems.

An inspection of the proof of both Theorems will reveal that the condition $t_n = n^{\delta/2}$ is only used at the last step to present the bound in its simplest form. A similar bound can be derived for other sequences t_n, for example, when t_n grows logarithmically, a fact which is used in the proof of Theorem 4.1 later.

3.3. Prior concentration for a class of GL priors

Proving concentration results for the GL family (2) in the general setting presents a much harder challenge compared to Theorem 3.3 since we now have

to additionally integrate over the n local parameters $\psi = (\psi_1, \ldots, \psi_n)$. We focus on an important sub-class in Theorem 3.4 below, namely the exponential Exp(λ) family for the distribution of g in (2). For analytical tractability, we additionally assume that θ_0 has only one non-zero entry. The interest in the Exp(λ) arises from the fact that normal-exponential scale mixtures give rise to the double-exponential family [25]: $\theta \mid \psi \sim N(0, \psi\sigma^2), \psi \sim \text{Exp}(1/2)$ implies $\theta \sim \text{DE}(\sigma)$, and hence this family of priors can be considered as a Bayesian version of the lasso [13]. We now state a concentration result for this class. The proof is provided in the supporting document.

Theorem 3.4. *Assume $\theta \sim$ GL with $g \in \mathcal{G}$ and $f \equiv \text{Exp}(\lambda)$ for some constant $\lambda > 0$. Also assume θ_0 has only one non-zero entry. Then, for a global constant $C_1 > 0$ depending only on M in the definition of $\mathcal{G}$, for all $\sqrt{w_n} = t_n < \|\theta_0\|_2 / 4$,*

$$\mathbb{P}(\|\theta - \theta_0\|_2 \leq t_n) \leq C_1 \int_{\psi_1=0}^{\infty} \frac{\psi_1^{(n-3)/2}}{\left\{\psi_1 + \|\theta_0\|_2^2 / (\pi w_n)\right\}^{(n-3)/2}} e^{-\psi_1} d\psi_1. \tag{9}$$

Let $v_n = r_n^2$ satisfy $v_n = O(\sqrt{n})$. Then, for $\|\theta_0\|_2 \geq 1/\sqrt{n}$,

$$\mathbb{P}(\|\theta - \theta_0\|_2 \leq r_n) \geq C_2 e^{-d_2\sqrt{n}} \int_{\psi_1=c_1\|\theta_0\|_2^2}^{\infty} \frac{\psi_1^{(n-3)/2}}{\left\{\psi_1 + \|\theta_0\|_2^2 / (\pi v_n)\right\}^{(n-3)/2}} e^{-\psi_1} d\psi_1, \tag{10}$$

where c_1, d_2, C_2 are positive global constants with $c_1 \geq 2$ and C_2 depends only on M in the definition of $\mathcal{G}$.

A more interpretable corollary can be stated as follows.

Corollary 3.5. *Assume $\theta \sim$ GL with $g \in \mathcal{G}$ and $f \equiv \text{Exp}(\lambda)$ for some constant $\lambda > 0$. Also assume θ_0 has only one non-zero entry and $\|\theta_0\|_2^2 > \log n$. Then, for a global constant $C > 0$ depending only on M in the definition of $\mathcal{G}$,*

$$\mathbb{P}(\|\theta - \theta_0\|_2 < \sqrt{\log n}) \leq e^{-C\sqrt{n}}. \tag{11}$$

Corollary 3.5 asserts that even in the simplest deviation from the null model with only one signal, one continues to have exponentially small concentration under an exponential prior on the local scales. From (5) in Theorem 3.1, appropriate point mass mixture priors (1) would have $\mathbb{P}(\|\theta - \theta_0\|_2 < t_n) \geq e^{-C\|\theta_0\|_1}$ under the same conditions as above, clearly showing that the wide difference in concentration still persists.

4. Posterior lower bounds in normal means

We have discussed the prior concentration for a high-dimensional vector θ without alluding to any specific model so far. In this section we show how prior concentration impacts posterior inference for the widely studied normal means problem[1] (see [26, 27, 10] and references therein):

$$y_i = \theta_i + \epsilon_i, \quad \epsilon_i \sim \mathrm{N}(0,1), \quad 1 \leq i \leq n. \tag{12}$$

The minimax rate s_n for the above model is given by $s_n^2 \asymp q_n \log(n/q_n)$ when $\theta_0 \in \ell_0[q_n; n]$. For this model [10] recently established that for point mass priors for θ with $\pi \sim \mathrm{beta}(1, \kappa n+1)$ and g_θ having Laplace like or heavier tails, the posterior contracts at the minimax rate, *i.e.*, $\mathbb{E}_{n,\theta_0}\mathbb{P}(\|\theta - \theta_0\|_2 < M s_n \mid y) \to 1$ for some constant $M > 0$. Thus we see that carefully chosen point mass priors are indeed optimal[2].

However not all choices for g_θ lead to optimal proceedures; [10] also showed that if g_θ is instead chosen to be standard Gaussian, *the posterior does not contract at the minimax rate*, *i.e.*, one could have $\mathbb{E}_{n,\theta_0}\mathbb{P}(\|\theta - \theta_0\|_2 < s_n \mid y) \to 0$ for signals of sufficiently large magnitude. To establish such a posterior lower-bound result, building on the work of [28], [10] showed that given a fixed sequence t_n, if there exists a sequence r_n ($r_n > t_n$) such that

$$\frac{\mathbb{P}(\|\theta - \theta_0\|_2 < t_n)}{\mathbb{P}(\|\theta - \theta_0\|_2 < r_n)} = o(e^{-r_n^2}), \tag{13}$$

then $\mathbb{P}(\|\theta - \theta_0\|_2 < t_n \mid y) \to 0$. This immediately shows the importance of studying the prior concentration. Intuitively, (13) would be satisfied when the prior mass of the bigger ball $\|\theta - \theta_0\|_2 < r_n$ is almost entirely contained in the annulus with inner radius t_n and outer radius r_n, so that the smaller ball $\|\theta - \theta_0\|_2 < t_n$ barely has any prior mass compared to the bigger ball. As an illustrative example, in the i.i.d. $\mathrm{N}(0,1)$ example with $t_n = s_n$, setting $r_n = \sqrt{n}$ would satisfy (13) above, proving that i.i.d. $\mathrm{N}(0,1)$ priors are sub-optimal. Our goal is to investigate whether a similar phenomenon persists for global-local priors in light of the concentration bounds developed in Theorems 3.3 and 3.4.

[1] Although we study the normal means problem, the ideas and results in this section are applicable to other models such as non-parametric regression and factor models.

[2] It is important that the hyper parameter for π depends on n. We do not know if the result holds without this

As in Section 3.2, we first state our posterior lower bound result for the case where there is only a global parameter.

Theorem 4.1. *Suppose we observe* $y \sim N_n(\theta_0, \mathrm{I}_n)$ *and* (12) *is fitted with a* GL *prior on* θ *such that* $\psi_j = 1$ *for all* j *and the prior* g *on the global parameter* τ *lies in* $\mathcal{G}$. *Assume* $\theta_0 \in \ell_0[q_n; n]$ *where* $q_n/n \to 0$ *and* $\|\theta_0\|_2 > s_n$, *with* $s_n^2 = q_n \log(n/q_n)$ *being the minimax squared error loss over* $\ell_0[q_n; n]$. *Then,* $\mathbb{E}_{n,\theta_0}\mathbb{P}(\|\theta - \theta_0\|_2 \leq \sqrt{A}s_n \mid y) \to 0$ *for any constant* $A > 0$.

Proof. Without loss of generality, assume $\|\theta_0\|_2 = o(\sqrt{n})$, since the posterior mass with a prior centered at the origin would be smaller otherwise. Choosing $t_n = \sqrt{A}s_n$, r_n to be a sequence such that $t_n < r_n < \|\theta_0\|_2$ and resorting to the two-sided bounds in Theorem 3.3, the ratio in (13) is smaller than $(t_n/r_n)^n$, and hence $e^{r_n^2}(t_n/r_n)^n \to 0$ since $r_n \leq \|\theta_0\|_2 = o(\sqrt{n})$. □

Theorem 4.1 states that a GL prior with only a global scale is sub-optimal if $\|\theta_0\|_2 > s_n$. Observe that in the complementary region $\{\|\theta_0\|_2 \leq s_n\}$, the estimator $\hat{\theta} \equiv 0$ attains squared error in the order of $q_n \log(n/q_n)$, implying the condition $\|\theta_0\|_2 > s_n$ is hardly stringent. Clearly, the absence of local scales makes it challenging to estimate both coefficients of different signal strengths simultaneously.

Next, we state a result for the sub-class of GL priors as in Theorem 3.4, i.e., when f has an exponential distribution leading to a double-exponential distribution marginally.

Theorem 4.2. *Suppose we observe* $y \sim N_n(\theta_0, \mathrm{I}_n)$ *and the model in* (12) *is fitted with a* GL *prior on* θ *such that* g *lies in* $\mathcal{G}$ *and* $f \equiv \mathrm{Exp}(\lambda)$ *for some constant* $\lambda > 0$. *Assume* $\theta_0 \in \ell_0[q_n; n]$ *with* $q_n = 1$ *and* $\|\theta_0\|_2^2 / \log n \to \infty$. *Then,* $\mathbb{E}_{n,\theta_0}\mathbb{P}(\|\theta - \theta_0\|_2 \leq \sqrt{A \log n} \mid y) \to 0$ *for any constant* $A > 0$.

From [10], appropriate point mass mixture priors would assign increasing mass with n to the neighborhood $\{\|\theta - \theta_0\|_2 \leq \sqrt{A \log n}\}$ for sufficiently large $A > 0$. Indeed, [26] established that the minimax rate for $\ell[1; n]$ is $\sqrt{2 \log n}\{1 + o(1)\}$. Hence, the Bayesian lasso [13] is sub-optimal even in the simplest deviation from the null model with only one moderately sized signal. We believe that the conclusions would continue to be valid if one only assumes f to have exponential tails plus some mild conditions on the behavior near zero. However, the assumptions of Theorem 4.2 precludes the case when f has polynomial tails, such as the horseshoe [16] and generalized double Pareto [17]. Very recently, [20] showed optimal posterior concentration for the

horseshoe prior where the global parameter τ is estimated using an empirical Bayes approach.

Another important question beyond the scope of the current paper should concern the behavior of the posterior when one plugs in an empirical Bayes estimator of the global parameter τ. However, we show below that the "optimal" sample-size dependent plug-in choice $\tau_n = c^2/\log n$ (so that marginally $\theta_j \sim \mathrm{DE}(c/\sqrt{\log n})$) for the lasso estimator [7] produces a sub-optimal posterior:

Theorem 4.3. *Suppose we observe $y \sim N_n(\theta_0, \mathrm{I}_n)$ and (12) is fitted with a* GL *prior on θ such that τ is deterministically chosen to be τ_n, i.e., $g \equiv \delta_{\tau_n}$ for a non-random sequence τ_n and $f \equiv \mathrm{Exp}(\lambda)$ for some constant $\lambda > 0$. Assume $\theta_0 \in \ell_0[q_n; n]$ with $q_n = 1$, $\|\theta_0\| > 4s_n$, $\|\theta_0\| \log n = o(n)$ and $\tau_n = c/\log n$ is used as the plug-in choice. Then, $\mathbb{E}_{n,\theta_0}\mathbb{P}(\|\theta - \theta_0\|_2 \leq s_n \mid y) \to 0$, with $s_n^2 = q_n \log(n/q_n)$ being the minimax squared error loss over $\ell_0[q_n; n]$.*

5. Auxiliary results

We collect a set of Lemmata we need to prove the main results. All proofs are deferred to the supporting document.

An important tool used throughout is a two-sided Gaussian small ball probability [29].

Lemma 5.1. *Suppose $\theta \sim \mathrm{N}_n(0, \Sigma)$ with Σ be a positive definite matrix and $\theta_0 \in \mathbb{R}^n$. Let $\|\theta_0\|_{\mathbb{H}}^2 = \theta_0^{\mathrm{T}}\Sigma^{-1}\theta_0$. Then, for any $t > 0$,*

$$e^{-\frac{1}{2}\|\theta_0\|_{\mathbb{H}}^2}\mathbb{P}(\|\theta\|_2 \leq t) \leq \mathbb{P}(\|\theta - \theta_0\|_2 \leq t). \tag{14}$$

In addition, if Σ is diagonal and $\|\theta_0\|_0 = 1$, i.e., θ_0 has only one non-zero entry, then for all $t < \frac{1}{4}\|\theta_0\|_2$,

$$\mathbb{P}(\|\theta - \theta_0\|_2 \leq t) \leq e^{-\frac{\|\theta_0\|_{\mathbb{H}}^2}{4}}\mathbb{P}(\|\theta\|_2 \leq t). \tag{15}$$

If the condition $\|\theta_0\|_0 = 1$ above is replaced by $\Sigma = \sigma^2 I$ for some $\sigma > 0$, then for all $t < \frac{1}{4}\|\theta_0\|_2$, the bound in (15) holds as well.

Remark 5.2. *It is well known that among balls of fixed radius, a zero mean multivariate normal distribution places the maximum mass on the ball centered at the origin. Lemma 5.1 provides a sharp bound on the probability of shifted balls in terms of the centered probability and the size of the shift, measured via the RKHS norm $\|\theta_0\|_{\mathbb{H}}^2$.*

We next state Lemma 5.3 to bound an incomplete gamma integral from below. Recall $\int_{\tau=0}^{\infty} \tau^{-n/2} e^{-a_n/(2\tau)} d\tau = \Gamma(n/2-1)(2/a_n)^{n/2-1}$. Lemma 5.3 shows that the same integral over $(0,1)$ is of the same order when $a_n \precsim n$.

Lemma 5.3. *For a sequence $a_n \leq n/(2e)$, $\int_{\tau=0}^{1} \tau^{-n/2} e^{-a_n/(2\tau)} d\tau \geq (2/a_n)^{n/2-1}\Gamma(n/2-1)\xi_n$, where $\xi_n \uparrow 1$ with $(1-\xi_n) \leq D/\sqrt{n}$ for some constant $D > 0$.*

We state a two sided bound for the complementary error function.

Lemma 5.4. *Let $erfc(x) = \frac{2}{\sqrt{\pi}} \int_x^{\infty} e^{-t^2} dt$ denote the complementary error function. Then,*

$$\sqrt{\pi} e^x \mathrm{erfc}(\sqrt{x}) \leq \frac{1}{\sqrt{x + 1/\pi}} \tag{16}$$

$$\sqrt{\pi} e^x \mathrm{erfc}(\sqrt{x}) \geq \left\{ \frac{1}{\sqrt{x}} \right\}^{1+\delta} \tag{17}$$

where (17) holds for any $\delta > 0$ provided $x \geq 2$.

We next state the Dirichlet integral formula (4.635 in [30]) to simplify a class of integrals over the simplex Δ^{n-1}:

Lemma 5.5. *Let $h(\cdot)$ be a Lebesgue integrable function and $\alpha_j > 0, j = 1, \ldots, n$. Then,*

$$\int_{\sum x_j \leq 1} h\big(\sum x_j\big) \prod_{j=1}^{n} x_j^{\alpha_j - 1} d\boldsymbol{x} = \frac{\prod_{j=1}^{n} \Gamma(\alpha_j)}{\Gamma\big(\sum_{j=1}^{n} \alpha_j\big)} \int_{t=0}^{1} h(t)\, t^{(\sum \alpha_j)-1} dt.$$

Lemma 5.5 follows simply by noting that the left hand side is $\mathbb{E}h(\sum_{j=1}^{n} X_j)$ up to normalizing constants where $(X_1, \ldots, X_n) \sim \mathrm{Diri}(\alpha_1, \ldots, \alpha_n, 1)$, so that $\sum_{j=1}^{n} X_j \sim \mathrm{Beta}(\sum \alpha_j, 1)$. Such probabilistic intuitions help us to reduce a more complicated integral stated in Lemma 5.6 below. The proof of Lemma 5.6 utilizes a beautiful identity found in [31]. We didn't find any reference for Lemma 5.6, though a related integral with $n/2$ in the exponent in the denominator appears in [30].

Lemma 5.6. *Let $q_j, j = 0, 1, \ldots, n$ be positive numbers. Then,*

$$\int_{\sum x_j \leq 1} \frac{\prod_{j=1}^{n} x_j^{-1/2}}{[\sum_{j=1}^{n} q_j x_j + q_0]^{n/2-1}} d\boldsymbol{x} = \frac{\Gamma(1/2)^n}{\Gamma(n/2)} q_0 (n/2-1) \int_{x=0}^{1} \frac{x^{n/2-2}(1-x)}{\prod_{j=1}^{n} \sqrt{q_j x + q_0}} dx.$$

6. Proofs of results in Sections 3 & 4

In this section, we prove the main results of the paper other than Theorem 3.4, which is deferred to the supporting document due to its length.

For GL shrinkage priors of the form (2), given $\boldsymbol{\psi} = (\psi_1, \ldots, \psi_n)^{\mathrm{T}}$ and τ, the elements of θ are conditionally independent with $\theta \mid \boldsymbol{\psi}, \tau \sim \mathrm{N}_n(0, \Sigma)$ with $\Sigma = \mathrm{diag}(\psi_1\tau, \ldots, \psi_n\tau)$. Hence we can use Lemma 5.1 to obtain

$$e^{-1/(2\tau)\sum_{j=1}^n \theta_{0j}^2/\psi_j}\, \mathbb{P}(\|\theta\|_2 < t_n \mid \boldsymbol{\psi}, \tau) \leq \mathbb{P}(\|\theta - \theta_0\|_2 \leq t_n \mid \boldsymbol{\psi}, \tau). \quad (18)$$

Furthemore, if $\psi_j = 1$ for all j, then again by using Lemma 5.1 we obtain

$$\mathbb{P}(\|\theta - \theta_0\|_2 \leq t_n \mid \boldsymbol{\psi}, \tau) \leq e^{-1/(4\tau)\sum_{j=1}^n \theta_{0j}^2} \mathbb{P}(\|\theta\|_2 < t_n \mid \boldsymbol{\psi}, \tau). \quad (19)$$

Letting $X_j = \theta_j^2$, X_j's are conditionally independent given $(\boldsymbol{\psi}, \tau)$ with X_j having a density $f(x_j \mid \boldsymbol{\psi}, \tau) = D/(\sqrt{\tau\psi_j x_j})e^{-x_j/(2\tau\psi_j)}$ on $(0, \infty)$, where $D = 1/(\sqrt{2\pi})$. Hence, with $w_n = t_n^2$,

$$\mathbb{P}(\|\theta\|_2 < t_n \mid \boldsymbol{\psi}, \tau) = D^n \int_{\sum x_j \leq w_n} \prod_{j=1}^n \frac{1}{\sqrt{x_j\tau\psi_j}} e^{-x_j/(2\tau\psi_j)} d\boldsymbol{x}. \quad (20)$$

For sake of brevity, we use $\{\sum x_j \leq w_n\}$ in (20) and all future references to denote the region $\{\boldsymbol{x} \in \mathbb{R}^n : x_j \geq 0\, \forall\, j = 1, \ldots, n, \sum_{j=1}^n x_j \leq w_n\}$. To estimate two-sided bounds for the marginal concentration $\mathbb{P}(\|\theta - \theta_0\|_2 \leq t_n)$, we need to combine (18) & (20) and integrate out $\boldsymbol{\psi}$ and τ carefully. We start by proving Theorem 3.2 & Theorem 3.3 where one only needs to integrate out τ.

Proof of Theorem 3.2

In (20), set $\psi_j = 1$ for all j, recall $D = 1/\sqrt{2\pi}$ and $w_n = t_n^2$, and integrate over τ to obtain,

$$\mathbb{P}(\|\theta\|_2 \leq t_n) = D^n \int_{\tau=0}^{\infty} g(\tau) \Big[\int_{\sum x_j \leq w_n} \prod_{j=1}^n \frac{1}{\sqrt{x_j\tau}} e^{-x_j/(2\tau)} d\boldsymbol{x} \Big] d\tau. \quad (21)$$

Substituting $g(\tau) = c\tau^{-(1+\alpha)}e^{-\beta/\tau}$ with $c = \beta^\alpha/\Gamma(\alpha)$ and using Fubini's theorem to interchange the order of integration between x and τ, (21) equals

$$\begin{aligned}
&cD^n \int_{\sum x_j \le w_n} \prod_{j=1}^n \frac{1}{\sqrt{x_j}} \left[\int_{\tau=0}^\infty \tau^{-(1+n/2+\alpha)} e^{-\frac{1}{2\tau}(2\beta+\sum x_j)} d\tau \right] d\boldsymbol{x} \\
&= cD^n 2^{n/2+\alpha}\Gamma(n/2+\alpha) \int_{\sum x_j \le w_n} \frac{1}{(2\beta+\sum x_j)^{n/2+\alpha}} \prod_{j=1}^n \frac{1}{\sqrt{x_j}}\, d\boldsymbol{x} \\
&= cD^n 2^{n/2+\alpha} w_n^{n/2}\Gamma(n/2+\alpha) \int_{\sum x_j \le 1} \frac{1}{(2\beta+w_n\sum x_j)^{n/2+\alpha}} \prod_{j=1}^n \frac{1}{\sqrt{x_j}}\, d\boldsymbol{x}.
\end{aligned} \tag{22}$$

Lemma 5.5 with $h(t) = 1/(2\beta + w_n t)^{n/2+\alpha}$ applied to (22) implies

$$\mathbb{P}(\|\theta\|_2 \le t_n) = cD^n 2^{n/2+\alpha} w_n^{n/2}\Gamma(n/2+\alpha)\frac{\Gamma(1/2)^n}{\Gamma(n/2)} \int_{t=0}^1 \frac{t^{n/2-1}}{(2\beta+w_n t)^{n/2+\alpha}} dt. \tag{23}$$

Substituting $D = 1/\sqrt{2\pi}$, bounding $(2\beta+w_n t)^{n/2+\alpha} \ge (2\beta)^{\alpha+1}(2\beta+w_n t)^{n/2-1}$, and letting $\tilde{w}_n = w_n/(2\beta)$, (23) can be bounded above by

$$\frac{\Gamma(n/2+\alpha)}{\Gamma(n/2)\Gamma(\alpha)(2\beta)^{\alpha+1}} \tilde{w}_n^{n/2} \int_{t=0}^1 \frac{t^{n/2-1}}{(1+\tilde{w}_n t)^{n/2-1}} dt \le \frac{w_n \Gamma(n/2+\alpha)}{\Gamma(n/2)\Gamma(\alpha)(2\beta)^{\alpha+1}} \left(\frac{\tilde{w}_n}{1+\tilde{w}_n} \right)^{n/2-1},$$

where the second inequality above uses $t/(a+t)$ is an increasing function in $t > 0$ for fixed $a > 0$. By definition, $w_n = n^\delta$ for $0 < \delta < 1$ and hence $\frac{w_n\Gamma(n/2+\alpha)}{\Gamma(n/2)\Gamma(\alpha)(2\beta)^{\alpha+1}}$ can be bounded above by $e^{C_1 \log n}$. Also, using $(1-x)^{1/x} \le e$ for all $x > 0$, $\{\tilde{w}_n/(1+\tilde{w}_n)\}^{n/2-1}$ can be bound above by $e^{-C_2 n/w_n} = e^{-C_2 n^{1-\delta}}$. Hence the overall bound is $e^{-Cn^{1-\delta}}$ for some appropriate constant $C > 0$. This completes the proof of Theorem 3.2. □

Proof of Theorem 3.3

We start with the upper bound in (7). The steps are similar as in the proof of Theorem 3.2 and hence only a sketch is provided. Set $w_n = t_n^2$. Bounding $g(\tau) \le M$ and interchanging order of integrals in (21),

$$\mathbb{P}(\|\theta\|_2 \le t_n) \le MD^n 2^{n/2-1}\Gamma(n/2-1) w_n \int_{\sum x_j \le 1} \frac{1}{(\sum x_j)^{n/2-1}} \prod_{j=1}^n \frac{1}{\sqrt{x_j}}\, d\boldsymbol{x}. \tag{24}$$

Invoking Lemma 5.5 with $h(t) = (1/t)^{n/2-1}$ in (26), we obtain

$$MD^n 2^{n/2-1}\Gamma(n/2-1)w_n \frac{\Gamma(1/2)^n}{\Gamma(n/2)} \int_{x=0}^{1} x^{n/2-1}/x^{n/2-1} dx = (M/2)\frac{w_n}{n/2-1} = C_2 n^{-(1-\delta)}. \tag{25}$$

From (26) and (25) we obtain

$$\mathbb{P}(\|\theta\|_2 \leq t_n) \leq C_2 n^{-(1-\delta)} \tag{26}$$

and thus the upper bound in (7) is proved.

We now turn towards proving the lower bound to the centered concentration in (7). Recalling that $g(\tau) \geq 1/M$ on $(0,1)$ for $g \in \mathcal{G}$, and interchanging integrals in (21), we have, with $K = 1/M$,

$$\mathbb{P}(\|\theta\|_2 \leq t_n) \geq KD^n \int_{\sum x_j \leq w_n} \prod_{j=1}^{n} \frac{1}{\sqrt{x_j}} \Big[\int_{\tau=0}^{1} \tau^{-n/2} e^{-\sum x_j/(2\tau)} d\tau \Big] d\boldsymbol{x}. \tag{27}$$

Clearly $\sum x_j \leq w_n$ and hence we can apply Lemma 5.3 in (27) to get

$$\mathbb{P}(\|\theta\|_2 \leq t_n) \geq K\xi_n D^n 2^{n/2-1}\Gamma(n/2-1)w_n \int_{\sum x_j \leq 1} \frac{1}{(\sum x_j)^{n/2-1}} \prod_{j=1}^{n} \frac{1}{\sqrt{x_j}}\, d\boldsymbol{x}. \tag{28}$$

The rest of the proof proceeds for showing the lower bound in (7) follows exactly as in the upper bound case from (26) onwards.

Finally, we combine (18) with (19) (with $\psi_j = 1$ for all j in this case) to bound the non-centered probability in (8). For the upper bound, we additionally use $g(\tau) \leq M$ for all τ to obtain

$$\mathbb{P}(\|\theta - \theta_0\|_2 \leq t_n) \leq MD^n \int_{\sum x_j \leq w_n} \prod_{j=1}^{n} \frac{1}{\sqrt{x_j}} \Big[\int_{\tau=0}^{\infty} \tau^{-n/2} e^{-[\frac{1}{2}\|\theta_0\|_2^2 + \sum x_j]/(2\tau)} d\tau \Big] d\boldsymbol{x} \tag{29}$$

$$= MD^n 2^{n/2-1}\Gamma(n/2-1)w_n^{n/2} \int_{\sum x_j \leq 1} \frac{1}{(\frac{1}{2}\|\theta_0\|_2^2 + w_n \sum x_j)^{n/2-1}} \prod_{j=1}^{n} \frac{1}{\sqrt{x_j}}\, d\boldsymbol{x} \tag{30}$$

$$= MD^n 2^{n/2-1}\Gamma(n/2-1)w_n^{n/2} \frac{\Gamma(1/2)^n}{\Gamma(n/2)} \int_{x=0}^{1} \frac{x^{n/2-1}}{(\frac{1}{2}\|\theta_0\|_2^2 + w_n x)^{n/2-1}} dx. \tag{31}$$

In the above display, (30) - (31) follows from applying Lemma 5.5 with $h(t) = 1/(\frac{1}{2}\|\theta_0\|_2^2 + w_n t)^{n/2-1}$. Simplifying constants in (31) as before and using $t/(a+t)$ is an increasing function in $t > 0$ for fixed $a > 0$, we complete the proof by bounding (31) above by

$$\frac{Cw_n}{(n/2-1)}\int_{x=0}^{1}\frac{(w_n x)^{n/2-1}}{(\frac{1}{2}\|\theta_0\|_2^2 + w_n x)^{n/2-1}}dx \leq \frac{Cw_n}{(n/2-1)}\left(\frac{w_n}{w_n + \frac{1}{2}\|\theta_0\|_2^2}\right)^{n/2-1} \leq \frac{Cw_n}{(n/2-1)}\left(\frac{w_n}{\frac{1}{2}\|\theta_0\|_2^2}\right)^{n/2-1}.$$

The right hand side of the above display can be bounded above by $e^{-cn\log a_n}$ for some constant $c > 0$.

For the lower bound on the prior concentration in the non-centered case, we combine (18) in the reverse direction along with (20). We then use the same idea as in the centered case to restrict the integral over τ to $(0,1)$ in (32). Thus we have

$$\mathbb{P}(\|\theta - \theta_0\|_2 \leq t_n) \geq KD^n \int_{\sum x_j \leq w_n}\prod_{j=1}^{n}\frac{1}{\sqrt{x_j}}\left[\int_{\tau=0}^{1}\tau^{-n/2}e^{-[\|\theta_0\|_2^2 + \sum x_j]/(2\tau)}d\tau\right]d\boldsymbol{x}. \tag{32}$$

Noting that $\|\theta_0\|_2^2 + \sum x_j \leq \|\theta_0\|_2^2 + w_n = o(n)$, we can invoke Lemma 5.3 to lower bound the inner integral over τ by $\xi_n\Gamma(n/2-1)2^{n/2-1}/(\|\theta_0\|_2^2 + \sum x_j)^{n/2-1}$ and proceed to obtain the same expressions as in (30) & (31) with M replaced by $K\xi_n$. The proof is then completed by observing that the resulting lower bound can be further bounded below as follows:

$$\frac{Cw_n}{(n/2-1)}\int_{x=0}^{1}\frac{(w_n x)^{n/2-1}}{(\|\theta_0\|_2^2 + w_n x)^{n/2-1}}dx \geq \frac{Cw_n}{(n/2-1)}\int_{x=1/2}^{1}\frac{(w_n x)^{n/2-1}}{(\|\theta_0\|_2^2 + w_n x)^{n/2-1}}dx$$
$$\geq \frac{Cw_n}{(n/2-1)}\left(\frac{w_n/2}{(\|\theta_0\|_2^2 + w_n/2)}\right)^{n/2-1} \geq \frac{Cw_n}{(n/2-1)}\left(\frac{w_n/2}{2\|\theta_0\|_2^2}\right)^{n/2-1},$$

where the last inequality uses $t_n \leq \|\theta_0\|_2$ so that $\|\theta_0\|_2^2 + w_n \leq 2\|\theta_0\|_2^2$. This finishes the proof of Theorem 3.3. □

Proof of Theorem 4.2

Let $m_n = (n-3)/2$. Fix $A > 0$. We set $t_n = s_n$, where $s_n = \sqrt{A\log n}$ so that $w_n = s_n^2 = A\log n$. Also, let $\|\theta_0\|_2^2 = 2\pi w_n u_n^2$, where u_n is a slowly

increasing sequence; we set $u_n = \log(\log n)$ for future references. Finally let $v_n = r_n^2 = \sqrt{m_n}$. With these choices, we proceed to show that (13) holds.

We first simplify (9) further. The function $x \to x/x(x+a)$ monotonically increases from 0 to 1 for any $a > 0$. Thus the integral in (9) can be further simplified as follows. For any $T_n > 0$,

$$\int_{\psi_1=0}^{\infty} \frac{\psi_1^{m_n}}{\left\{\psi_1 + \|\theta_0\|_2^2/(2\pi w_n)\right\}^{m_n}} e^{-\psi_1} d\psi_1$$
$$\leq \int_{\psi_1=0}^{T_n} \frac{\psi_1^{m_n}}{\left\{\psi_1 + \|\theta_0\|_2^2/(2\pi w_n)\right\}^{m_n}} e^{-\psi_1} d\psi_1 + \int_{\psi_1=T_n}^{\infty} e^{-\psi_1} d\psi_1 \leq \left(\frac{T_n}{T_n + u_n^2}\right)^{m_n} + e^{-T_n}. \tag{33}$$

We choose an appropriate T_n which gives us the necessary bound, namely $T_n = u_n\sqrt{m_n}$. Then, using the fact that $(1-x)^{1/x} \leq e^{-1}$ for all $x \in (0,1)$, we have

$$\left(\frac{T_n}{T_n + u_n^2}\right)^{m_n} = \left(\frac{\sqrt{m_n}}{\sqrt{m_n} + u_n}\right)^{m_n} = \left(1 - \frac{u_n}{\sqrt{m_n} + u_n}\right)^{m_n} \leq e^{-m_n u_n/(\sqrt{m_n}+u_n)} \leq e^{-u_n\sqrt{m_n}/2},$$

where for the last part used that $e^{-1/x}$ is an increasing function and $\sqrt{m_n} + u_n \leq 2\sqrt{m_n}$. Thus, using (9) and substituting T_n in (33) yields, for a global constant $C_1 > 0$,

$$\mathbb{P}(\|\theta - \theta_0\|_2 \leq s_n) \leq \frac{C_1 w_n}{(n/2-1)} \sqrt{\frac{w_n}{\|\theta_0\|_2^2}}\, e^{-u_n\sqrt{m_n}/2}. \tag{34}$$

Next, again using the fact that $x \to x/x(x+a)$ is monotonically increasing, and choosing $c_2 = \infty$, we simplify the lower bound (10). We have

$$\int_{\psi_1=c_1\|\theta_0\|^2}^{\infty} \frac{\psi_1^{m_n}}{\left\{\psi_1 + \|\theta_0\|_2^2/(\pi v_n)\right\}^{m_n}} e^{-\psi_1} d\psi_1 \;\; \geq \left(\frac{v_n}{v_n + C}\right)^{m_n} e^{-c_1\|\theta_0\|_2^2}, \tag{35}$$

for some constant $C > 0$. Since $(1-x)^{1/x} \geq e^{-2}$ for all $x \in (0, 1/2)$ and $e^{-1/x}$ is an increasing function in $x > 0$, we have,

$$\left(\frac{v_n}{v_n + C}\right)^{m_n} \geq e^{-\sqrt{m_n}/2}.$$

Hence, the left hand side of (35) is bounded below by $e^{-(\sqrt{m_n}+2c_1\|\theta_0\|_2^2)/2}$, resulting in

$$\mathbb{P}(\|\theta - \theta_0\|_2 \leq r_n) \geq \frac{C_2\xi_n v_n}{(n/2-1)} \sqrt{\frac{v_n}{v_n + \|\theta_0\|_2^2}}\, e^{-(\sqrt{m_n}+2c_1\|\theta_0\|_2^2)/2}. \tag{36}$$

Thus, finally, noting that $u_n \to \infty$, (13) follows since

$$\frac{\mathbb{P}(||\theta-\theta_0||_2 < s_n)}{\mathbb{P}(||\theta-\theta_0||_2 < r_n)} \times e^{r_n^2} \leq C' \frac{w_n^{3/2}}{v_n} e^{C(\sqrt{m_n}+\sqrt{n}+\|\theta_0\|_2^2)}\, e^{-u_n\sqrt{m_n}/2} \to 0,$$

where $C, C' > 0$ are constants. This finishes the proof of Theorem 4.2. □

Proof of Theorem 4.3

As before, we assume $\lambda = 1$ without loss of generality, since it can be absorbed in the constant appearing the sequence τ_n otherwise. For $t_n < \frac{1}{4}\|\theta_0\|_2 = \frac{1}{4}|\theta_{01}|$, using (15) we obtain

$$\mathbb{P}(\|\theta-\theta_0\| < t_n) \leq D^n \tau_n^{-n/2} w_n^{n/2} \int_{\sum x_j \leq 1} \prod_{j=1}^n \frac{1}{\sqrt{x_j}} \Big\{ \prod_{j=1}^n \int_{\psi_j=0}^\infty \frac{e^{-\psi_j}}{\sqrt{\psi_j}} \exp\Big(- \frac{w_n x_j + \theta_{01}^2/2}{2\tau_n \psi_j}\Big) d\psi_j \Big\} dx,$$

where $w_n = t_n^2$. Using the fact $\int_0^\infty \frac{1}{\sqrt{x}} \exp\big\{ - \big(\frac{a}{x} + x\big)\big\} dx = \sqrt{\pi} e^{-2\sqrt{a}}$, we obtain

$$\begin{aligned}
&\mathbb{P}(\|\theta-\theta_0\| < t_n) \\
&\leq D^n \pi^{n/2} \tau_n^{-n/2} w_n^{n/2} \int_{\sum x_j \leq 1} \prod_{j=1}^n \frac{1}{\sqrt{x_j}} \exp\Big\{ - 2\sqrt{\frac{w_n x_j + \theta_{01}^2/2}{2\tau_n}}\Big\} dx \\
&\leq \Big(\frac{D^2 \pi w_n}{\tau_n}\Big)^{n/2} e^{-\frac{|\theta_{01}|}{\sqrt{\tau_n}}} \int_{\sum x_j \leq 1} \prod_{j=1}^n \frac{1}{\sqrt{x_j}} dx = \Big(\frac{D^2 \pi w_n}{\tau_n}\Big)^{n/2} e^{-\frac{|\theta_{01}|}{\sqrt{\tau_n}}} \frac{\Gamma(1/2)^n}{\Gamma(n/2+1)},
\end{aligned} \tag{37}$$

where the second to third inequality uses $x_j \geq 0$ and the last integral follows from Lemma 5.5. Along the same lines,

$$\begin{aligned}
\mathbb{P}(\|\theta-\theta_0\| < r_n) &\geq \Big(\frac{D^2 \pi v_n}{\tau_n}\Big)^{n/2} \int_{\sum x_j \leq 1} \prod_{j=1}^n \frac{1}{\sqrt{x_j}} \exp\Big\{ - 2\sqrt{\frac{v_n x_j + \theta_{01}^2}{2\tau_n}}\Big\} dx \\
&\geq \Big(\frac{D^2 \pi v_n}{\tau_n}\Big)^{n/2} e^{-\frac{\sqrt{2}(|\theta_{01}|+\sqrt{nv_n})}{\sqrt{\tau_n}}} \frac{\Gamma(1/2)^n}{\Gamma(n/2+1)},
\end{aligned} \tag{38}$$

where $v_n = r_n^2/4$. From the second to third equation in the above display, we used $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ and $\sum_{j=1}^n \sqrt{x_j} \leq \sqrt{n}$ by Cauchy-Schwartz inequality if $x \in \Delta^{n-1}$. Thus, from (37) & (38), the ratio in (13) can be bounded above as:

$$\frac{\mathbb{P}(\|\theta-\theta_0\| < t_n)}{\mathbb{P}(\|\theta-\theta_0\| < r_n)} \leq \Big(\frac{w_n}{v_n}\Big)^{n/2} e^{\frac{\sqrt{2v_n n} + (\sqrt{2}-1)|\theta_{01}|}{\tau_n}}.$$

Choose $t_n = s_n$, $r_n = 2\sqrt{2}s_n$ so that $v_n = 2w_n = 2q_n \log(n/q_n)$ and $(w_n/v_n)^{n/2} = e^{-Cn}$. Clearly $v_n n/\tau_n \leq Cnq_n(\log n)^2$ and hence, $e^{\sqrt{2v_n n/\tau_n}} = o(e^{Cn})$ by assumption. Also, $|\theta_{01}|/\tau_n = o(n)$. Thus, the right hand side of the above display $\to 0$, proving the assertion of the Theorem.

Acknowledgements

NSP is indebted to Prof. Gine for generously sharing his time and ideas. We thank Professors Sara van de Geer and Richard Nickl for encouragement and the referees for constructive criticism.

References

[1] P. Bühlmann, S. Van De Geer, Statistics for high-dimensional data: methods, theory and applications, Springer, 2011.

[2] P. Zhao, B. Yu, On model selection consistency of lasso, Journal of Machine Learning Research 7 (2) (2006) 2541–2567.

[3] S. Van de Geer, High-dimensional generalized linear models and the lasso, The Annals of Statistics 36 (2) (2008) 614–645.

[4] C. Zhang, J. Huang, Model selection consistency of the lasso selection in high-dimensional linear regression, The Annals of Statistics 36 (4) (2008) 1567–1594.

[5] N. Meinshausen, B. Yu, Lasso-type recovery of sparse representations for high-dimensional data, The Annals of Statistics 37 (1) (2009) 246–270.

[6] G. Raskutti, M. Wainwright, B. Yu, Minimax rates of estimation for high-dimensional linear regression over l_q balls, Information Theory, IEEE Transactions on 57 (10) (2011) 6976–6994.

[7] S. Negahban, P. Ravikumar, M. Wainwright, B. Yu, A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers, arXiv preprint arXiv:1010.2731.

[8] B. Efron, T. Hastie, I. Johnstone, R. Tibshirani, Least angle regression, The Annals of statistics 32 (2) (2004) 407–499.

[9] R. Nickl, S. van de Geer, et al., Confidence sets in sparse regression, The Annals of Statistics 41 (6) (2013) 2852–2876.

[10] I. Castillo, A. van der Vaart, Needles and straws in a haystack: Posterior concentration for possibly sparse sequences, The Annals of Statistics 40 (4) (2012) 2069–2101.

[11] I. Castillo, J. Schmidt-Hieber, A. W. van der Vaart, Bayesian linear regression with sparse priors, arXiv preprint arXiv:1403.0735.

[12] N. Polson, J. Scott, Shrink globally, act locally: Sparse Bayesian regularization and prediction, in: Bayesian Statistics 9 (J.M. Bernardo, M.J. Bayarri, J.O. Berger, A.P. Dawid, D. Heckerman, A.F.M. Smith and M. West, eds.), Oxford University Press, New York, 2010, pp. 501–538.

[13] T. Park, G. Casella, The bayesian lasso, Journal of the American Statistical Association 103 (482) (2008) 681–686.

[14] M. Tipping, Sparse bayesian learning and the relevance vector machine, The Journal of Machine Learning Research 1 (2001) 211–244.

[15] J. Griffin, P. Brown, Inference with normal-gamma prior distributions in regression problems, Bayesian Analysis 5 (1) (2010) 171–188.

[16] C. Carvalho, N. Polson, J. Scott, The horseshoe estimator for sparse signals, Biometrika 97 (2) (2010) 465–480.

[17] A. Armagan, D. B. Dunson, J. Lee, Generalized double pareto shrinkage, Statistica Sinica 23 (1) (2013) 119.

[18] D. Pati, A. Bhattacharya, N. S. Pillai, D. Dunson, et al., Posterior contraction in sparse bayesian factor models for massive covariance matrices, The Annals of Statistics 42 (3) (2014) 1102–1130.

[19] A. Bhattacharya, D. Pati, N. S. Pillai, D. B. Dunson, Dirichlet-laplace priors for optimal shrinkage, Journal of the American Statistical Association To appear.

[20] S. van der Pas, B. Kleijn, A. van der Vaart, et al., The horseshoe estimator: Posterior concentration around nearly black vectors, Electronic Journal of Statistics 8 (2) (2014) 2585–2618.

[21] J. Scott, J. Berger, Bayes and empirical-bayes multiplicity adjustment in the variable-selection problem, The Annals of Statistics 38 (5) (2010) 2587–2619.

[22] C. Hans, Elastic net regression modeling with the orthant normal prior, Journal of the American Statistical Association 106 (496) (2011) 1383–1393.

[23] A. Gelman, Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper), Bayesian analysis 1 (3) (2006) 515–534.

[24] N. Polson, J. Scott, On the half-cauchy prior for a global scale parameter, arXiv preprint arXiv:1104.4937.

[25] M. West, On scale mixtures of normal distributions, Biometrika 74 (3) (1987) 646–648.

[26] D. Donoho, I. Johnstone, J. Hoch, A. Stern, Maximum entropy and the nearly black object, Journal of the Royal Statistical Society. Series B (Methodological) (1992) 41–81.

[27] I. Johnstone, B. Silverman, Needles and straw in haystacks: Empirical bayes estimates of possibly sparse sequences, The Annals of Statistics 32 (4) (2004) 1594–1649.

[28] I. Castillo, et al., Lower bounds for posterior rates with gaussian process priors, Electronic Journal of Statistics 2 (2008) 1281–1299.

[29] A. van der Vaart, J. van Zanten, Reproducing kernel Hilbert spaces of Gaussian priors, IMS Collections 3 (2008) 200–222.

[30] I. Gradshteyn, I. Ryzhik, Corrected and enlarged edition, Tables of Integrals, Series and ProductsAcademic Press, New York.

[31] J. Dickey, Three multidimensional-integral identities with bayesian applications, The Annals of Mathematical Statistics 39 (5) (1968) 1615–1628.